



March 6, 2023

National Telecommunications and Information Administration
1401 Constitution Avenue, NW
Washington, DC 20230

Submitted via Regulations.gov

Re: Privacy, Equity, and Civil Rights Request for Comment (NTIA-2023-0001)

Upturn is a non-profit organization that advances equity and justice in the design, governance, and use of technology. We submit this comment in response to the National Telecommunications and Information Administration’s (NTIA) Privacy, Equity and Civil Rights Request for Comment. We hope that this comment will help to inform the NTIA’s forthcoming “report on whether and how commercial data practices can lead to disparate impacts and outcomes for marginalized or disadvantaged communities.”¹

Our comment makes five key points.

1. For nearly a decade, the federal government has documented how automated systems can create and exacerbate discrimination.
2. As the federal government moves toward addressing algorithmic discrimination, federal agencies need new resources, structures, and authorities.
3. The federal government should provide clear guidance to companies about collecting or inferring sensitive demographic data for anti-discrimination testing purposes.
4. The federal government should provide clear guidance to companies about measuring algorithmic discrimination and searching for less discriminatory alternative models.
5. The federal government should significantly expand its anti-discrimination testing capabilities to uncover algorithmic discrimination.

¹ National Telecommunications and Information Administration, Department of Commerce, Privacy, Equity, and Civil Rights Request for Comment, 88 FR 3714, January 20, 2023, *available at* <https://www.federalregister.gov/documents/2023/01/20/2023-01088/privacy-equity-and-civil-rights-request-for-comment>.

1. For nearly a decade, the federal government has documented how automated systems can create and exacerbate discrimination.

It is well documented that powerful institutions now use a variety of automated, data-driven technologies to shape key decisions about people’s lives. These technologies can both create and exacerbate racial and economic disparities in housing, employment, public benefits, education, the criminal legal system, healthcare, and other areas of opportunity and wellbeing. A series of federal government reports, white papers, documents, and other requests for information or comments have clearly highlighted these problems.

In 2014, during the Obama administration, the White House released a landmark report titled “Big Data: Seizing Opportunities, Preserving Values,”² which the NTIA helped to draft. The report found that “big data technologies can cause societal harms beyond damages to privacy, such as discrimination against individuals and groups. This discrimination can be the inadvertent outcome of the way big data technologies are structured and used. It can also be the result of intent to prey on vulnerable classes.”³

A 2016 report from the Federal Trade Commission (FTC), “Big Data A Tool for Inclusion or Exclusion?”, sought to “educate businesses on important laws and research that are relevant to big data analytics and provide suggestions aimed at maximizing the benefits and minimizing its risks.”⁴ Later in 2016, the White House’s “Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights” further recognized that “if these technologies are not implemented with care, they can also perpetuate, exacerbate, or mask harmful discrimination.”⁵

During the Biden-Harris administration, in August 2022, the FTC issued an Advance Notice of Proposed Rulemaking (ANPR) on Commercial Surveillance and Data

² White House, *Big Data: Seizing Opportunities, Preserving Values* (May 2014), available at https://obamawhitehouse.archives.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf.

³ *Id.*, at 51.

⁴ Federal Trade Commission, *Big Data: A Tool for Inclusion or Exclusion? Understanding the Issues* at 1 (Jan. 2016), available at <https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf>.

⁵ White House, *Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights*, May 2016, 5, available at https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/2016_0504_data_discrimination.pdf.

Security.⁶ Among other things, the FTC’s ANPR sought comment on how the Commission should address algorithmic discrimination. The Commission received hundreds of comments regarding these questions, including from the NTIA, that demonstrate the pervasiveness of discriminatory commercial surveillance practices, both historically and today.⁷

In October 2022, the White House Office of Science and Technology Policy published a “Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People.”⁸ The Blueprint not only documented the many ways that automated systems create and exacerbate, discrimination, but also detailed the types of steps that companies could take to detect, prevent, and remediate discrimination.

Finally, in February 2023, President Biden signed an Executive Order on racial equity that directs agencies to “comprehensively use their respective civil rights authorities and offices to prevent and address discrimination and advance equity for all,” which includes “protecting the public from algorithmic discrimination.”⁹ That same Executive Order says that when “designing, developing, acquiring, and using artificial intelligence and automated systems in the Federal Government, agencies shall do so, consistent with applicable law, in a manner that advances equity.”¹⁰ This Executive Order appears to be the first time that an administration has formally defined “algorithmic discrimination.”

⁶ Federal Trade Commission, *Trade Regulation Rule on Commercial Surveillance and Data Security*, 87 FR 51273 (Aug. 22, 2022), available at <https://www.federalregister.gov/documents/2022/08/22/2022-17752/trade-regulation-rule-on-commercial-surveillance-and-data-security>.

⁷ See, e.g., organizational comments submitted in response to the FTC ANPR on Commercial Surveillance and Data Security, 87 FR 51273 by Upturn (Nov. 21, 2022), available at <https://www.upturn.org/static/files/Upturn-FTC-ANPR-Comment-20221121.pdf>; by Lawyers’ Committee for Civil Rights Under Law (Nov. 21, 2022), available at <https://www.regulations.gov/comment/FTC-2022-0053-0982>; by NAACP Legal Defense and Educational Fund, Inc. (Nov. 21, 2022), available at <https://www.naacpldf.org/wp-content/uploads/LDF-Comment-on-FTC-Rulemaking-on-Commercial-Surveillance-and-Data-Security27.pdf>; and by the American Civil Liberties Union (Nov. 21, 2022), available at <https://www.aclu.org/letter/aclu-comment-commercial-surveillance-anpr-r111004>.

⁸ White House Office of Science and Technology Policy, *Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People* at 3 (Oct. 4 2022), available at <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>.

⁹ White House, *Executive Order on Further Advancing Racial Equity and Support for Underserved Communities Through The Federal Government* (Feb. 16, 2023), available at <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/02/16/executive-order-on-further-advancing-racial-equity-and-support-for-underserved-communities-through-the-federal-government/>.

¹⁰ *Id.*

These are just a few examples of efforts by the federal government that have documented the many ways that automated systems can exacerbate structural inequities in our society. Particularly in response to Question 2 regarding specific examples of the disproportionate negative outcomes that marginalized communities experience, the NTIA should consult these past efforts.

2. As the federal government moves toward addressing algorithmic discrimination, federal agencies need new resources, structures, and authorities.

While the term “algorithmic discrimination” may be relatively new, the technologies, practices, and harms in question often are not. As detailed in many of the reports above, the algorithmic and other data-driven technologies that exacerbate racial, gender, disability, and other forms of discrimination were often developed decades ago. An earlier generation of statistical models preceded the more complex tools that rely on machine learning and other newer techniques in use today. But even as techniques evolve, the underlying problems and material harms remain the same and continue to this day. Though technologies new and old routinely mediate access to opportunity in traditionally covered civil rights areas like housing, employment, and credit, longstanding civil rights protections and antidiscrimination laws have not kept pace with technological change.

Since the beginning of the Biden-Harris administration, civil rights and technology justice groups, including Upturn, have urged the White House to center civil rights and racial equity in its artificial intelligence and technology policy priorities, and have offered clear recommendations on how the administration can address technology’s role in hiring, housing, and financial services discrimination.¹¹

Several federal agencies have recently taken concrete actions to ensure that existing authorities, regulations, and guidelines are understood to cover algorithmic discrimination, where applicable. For example:

¹¹ See Upturn, ACLU, the Leadership Conference on Civil and Human Rights, et al., *Memos on Centering Civil Rights in Artificial Intelligence and Technology Policy* (Jul. 13, 2021), available at <https://www.upturn.org/static/files/2021-07-13%20Coalition%20Letter%20to%20OSTP%20on%20Centering%20Civil%20Rights%20in%20AI%20Policy.pdf>; see also Leadership Conference on Civil and Human Rights, et al., *Letter to Ambassador Rice on Civil Rights and AI* (Oct. 27, 2021), available at <https://civilrights.org/resource/letter-to-ambassador-rice-on-civil-rights-and-ai>.

- The Department of Justice and Equal Employment Opportunity Commission released updated guidance and technical assistance documents clarifying how new hiring technologies can lead to discrimination in violation of the Americans with Disabilities Act.¹²
- A forthcoming rule by the Federal Housing Finance Agency, Federal Reserve Board, Office of the Comptroller of the Currency, Federal Deposit Insurance Corporation, National Credit Union Administration, and Consumer Financial Protection Bureau regarding use of automated valuation models will address potential bias in mortgage appraisals by including a nondiscrimination quality control standard in the proposed rule.¹³
- The Department of Housing and Urban Development will be issuing guidance on tenant screening algorithms and their use under the Fair Housing Act.¹⁴
- The Department of Justice filed a Statement of Interest in a case, clarifying that the Fair Housing Act applies to tenant screening companies.¹⁵
- The Department of Health and Human Services has a proposed rule that would prohibit clinical algorithms from discriminating against individuals based on race, color, national origin, sex, age, or disability.¹⁶

¹² Department of Justice, Civil Rights Division, *Algorithms, Artificial Intelligence, and Disability Discrimination in Hiring* (May 12, 2022), <https://www.ada.gov/resources/ai-guidance>; Equal Employment Opportunity Commission, *The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees* (May 12, 2022), available at <https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificial-intelligence>.

¹³ Interagency Task Force on Property Appraisal and Valuation Equity, *Action Plan to Advance Property Appraisal and Valuation Equity Closing the Racial Wealth Gap by Addressing Mis-valuations for Families and Communities of Color* (Mar. 2022), available at <https://pave.hud.gov/sites/pave.hud.gov/files/documents/PAVEActionPlan.pdf>.

¹⁴ White House, *Fact Sheet: Biden-Harris Administration Announces Key Actions to Advance tech Accountability and Protect the Rights of the American Public*, (Oct. 4, 2022), available at <https://www.whitehouse.gov/ostp/news-updates/2022/10/04/fact-sheet-biden-harris-administration-announces-key-actions-to-advance-tech-accountability-and-protect-the-rights-of-the-american-public>.

¹⁵ Statement of Interest of the United States, *Louis et al. v. SafeRent et al.*, D. Mass. 22-cv-10800 (Jan. 9, 2023), available at <https://www.justice.gov/opa/press-release/file/1561526/download>.

¹⁶ Centers for Medicare and Medicaid Services, *Notice of Proposed Rulemaking re: Nondiscrimination in Health Programs and Activities*, 87 FR 47824 (Aug. 4, 2022), available at <https://www.federalregister.gov/documents/2022/08/04/2022-16217/nondiscrimination-in-health-programs-and-activities>.

- The Equal Employment Opportunity Commission’s draft strategic enforcement plan squarely focuses on the use of algorithmic systems throughout the hiring process.¹⁷
- The General Counsel for the National Labor Relations Board issued a memo noting that workplace surveillance (from intrusive or abusive electronic monitoring to automated management practices) significantly harms employees’ ability to engage in protected activity such as unionizing.¹⁸
- The Consumer Financial Protection Bureau clarified that ECOA and Regulation B do not permit creditors to use complex algorithms when doing so means they cannot provide the specific and accurate reasons for adverse actions.¹⁹
- The Department of Justice settlement with Meta regarding alleged violations of the Fair Housing Act required the company to develop the Variance Reduction System, which aims to reduce the demographic variance between an eligible ad audience and the actual audience that is shown a housing ad.²⁰

These agency efforts are clear steps in the right direction, but more is required. The NTIA should build on the directives of the February 2023 Executive Order on racial equity and help to ensure that the federal agencies have the resources they need to “prevent and remedy discrimination, including by protecting the public from algorithmic discrimination.”²¹ These resources include the capacity and expertise to perform discrimination testing when it comes to automated systems, as described in Point 5 below. The NTIA should also create or formalize the government structures necessary to enable robust interagency collaboration on shared regulatory priorities on these issues. And the NTIA should conduct a comprehensive examination of our nation’s civil rights laws, to

¹⁷ Equal Employment Opportunity Commission, *Request for information and comment re: Draft Strategic Enforcement Plan*, 88 FR 1379 (Jan. 10, 2023), available at <https://www.federalregister.gov/documents/2023/01/10/2023-00283/draft-strategic-enforcement-plan>.

¹⁸ National Labor Relations Board, *NLRB General Counsel Issues Memo on Unlawful Electronic Surveillance and Automated Management Practices* (Oct. 31, 2022), <https://www.nlr.gov/news-outreach/news-story/nlr-general-counsel-issues-memo-on-unlawful-electronic-surveillance-and>.

¹⁹ Consumer Financial Protection Bureau, *Consumer Financial Protection Circular 2022-03* (May 26, 2022), available at <https://www.consumerfinance.gov/compliance/circulars/circular-2022-03-adverse-action-notification-requirements-in-connection-with-credit-decisions-based-on-complex-algorithms>.

²⁰ Department of Justice, *Justice Department Secures Groundbreaking Settlement Agreement with Meta Platforms, Formerly Known as Facebook, to Resolve Allegations of Discriminatory Advertising* (Jun. 21, 2022), <https://www.justice.gov/opa/pr/justice-department-secures-groundbreaking-settlement-agreement-meta-platforms-formerly-known>.

²¹ White House, *supra* note 9.

suggest updates that are necessary to clarify and strengthen the authorities that federal agencies have to address algorithmic discrimination.

3. The federal government should provide clear guidance to companies about collecting or inferring sensitive demographic data for anti-discrimination testing purposes.

When companies discriminate, whether intentionally or not, people can be unfairly hampered in their pursuit of basic services and economic opportunities, such as stable housing, quality jobs, and financial security. Large companies operating in traditionally covered civil rights areas should be expected to perform anti-discrimination testing when they build or use automated systems. To do so, companies will need some method of gathering basic demographic data about their users, applicants, or consumers.

The NTIA should explicitly call for a reconsideration of existing legal and regulatory prohibitions on the collection of demographic data in civil rights contexts and advocate for affirmative obligations, given the importance of this data for anti-discrimination testing. It should also clearly state the importance of requirements that demographic data be stored separately from other data, and should only be accessible for anti-discrimination purposes. The NTIA should help the administration ensure that ongoing legislative and regulatory data privacy efforts do not interfere with the collection and use of sensitive demographic data for anti-discrimination purposes.

Testing algorithms for discrimination generally requires covered entities to either collect or infer demographic data. In many cases, entities may not already have access to such data, either by choice or because of legal prohibition. When entities that don't have relevant demographic data try to measure disparities, they often turn to proxies based on available information. For example, methodologies such as Bayesian Improved Surname Geocoding (BISG) or Bayesian Improved First Name Surname Geocoding (BIFSG) use people's names and geographic locations to infer race and ethnicity. A "growing number of companies are turning to inferential methods [such as BISG or BIFSG] in their efforts to measure discrimination, even in the absence of clear legal or organizational guidance."²²

²² Aaron Rieke, Vincent Southerland, Dan Svirsky, Mingwei Hsu, *Imperfect Inferences: A Practical Assessment*, Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (2022), available at <https://dl.acm.org/doi/fullHtml/10.1145/3531146.3533140>.

As an illustrative example, lenders are generally prohibited under the Equal Credit Opportunity Act and Regulation B from collecting information on consumers' race and ethnicity for non-mortgage credit products. Exceptions include applications for home mortgages covered under the Home Mortgage Disclosure Act and for applications for credit for women-owned, minority-owned, and small businesses.²³ However, information about consumers' race and ethnicity is required for fair lending purposes. As a result, the CFPB's Office of Research and Division of Supervision, Enforcement, and Fair Lending relies "on a BISG proxy probability for race and ethnicity in fair lending analysis conducted for non-mortgage products."²⁴ The FTC also relies on BISG/BIFSG in their research efforts.²⁵

Inference methodologies such as BISG/BIFSG can "have real, practical utility in shedding light on aggregate, directional disparities."²⁶ Use of inference methodologies like BISG is what empowers researchers to measure racial disparities in tax audits,²⁷ allows the monitorship of financial technology company Upstart to document racial disparities and search for less discriminatory alternatives,²⁸ and makes DOJ's fair housing settlement with Meta possible.²⁹ Nevertheless, there are known limitations. For example, BISG appears to

²³ 12 U.S.C. 2803 (a)(1); 15 U.S.C. 1691c-2.

²⁴ Consumer Financial Protection Bureau, *Using publicly available information to proxy for unidentified race and ethnicity: A methodology and assessment* at 23 (2014), available at https://files.consumerfinance.gov/f/201409_cfpb_report_proxy-methodology.pdf.

²⁵ Federal Trade Commission, *Protecting Older Consumers 2021-2022: A Report of the Federal Trade Commission* at 39 (Oct. 18, 2022), available at https://www.ftc.gov/system/files/ftc_gov/pdf/P144400OlderConsumersReportFY22.pdf.

²⁶ Aaron Rieke, Vincent Southerland, Dan Svirsky, Mingwei Hsu, *Imperfect Inferences: A Practical Assessment*, Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (2022), available at <https://dl.acm.org/doi/fullHtml/10.1145/3531146.3533140>.

²⁷ Hadi Elzayn, Evelyn Smith, Thomas Hertz, Arun Ramesh, Robin Fisher, Daniel E. Ho, Jacob Goldin, *Measuring and Mitigating Racial Disparities in Tax Audits* (Jan. 30, 2023), available at <https://siepr.stanford.edu/publications/measuring-and-mitigating-racial-disparities-tax-audits>.

²⁸ Relman Colfax, *Fair Lending Monitorship of Upstart Network's Lending Model*, <https://www.relmanlaw.com/cases-406>.

²⁹ Department of Justice, *Justice Department and Meta Platforms Inc. Reach Key Agreement as They Implement Groundbreaking Resolution to Address Discriminatory Delivery of Housing Advertisements* (Jan. 9, 2023), <https://www.justice.gov/opa/pr/justice-department-and-meta-platforms-inc-reach-key-agreement-they-implement-groundbreaking>.

be more accurate for older populations,³⁰ and can sometimes understate or overstate underlying disparities.³¹

Other inference methodologies rely on slightly different techniques. For example, to measure discrimination on its platform, Airbnb’s Project Lighthouse relies on class labels inferred by a third-party partner, which are subject to a strict privacy and security protocol.³² That third-party partner is provided with “pairs of first names and photos of Airbnb hosts or guests,” and “then assign[s] a perceived race to each pair of first names and photos . . . the Research Partner [is required] to engage in the human perception of race, i.e. to not utilize algorithms for computer perception of race.”³³ Airbnb argues because discrimination is based on perception, an inference methodology that replicates perception makes sense for its measurement purposes.³⁴

In addition to racial inferences, efforts such as the Federal Equitable Data Working Group are working to “establish best practices for measuring sexual orientation, gender identity, disability, and rural location.”³⁵

³⁰ Dzifa Adjaye-Gbewonyo, Robert A. Bednarczyk, Robert L. Davis, Saad B. Omer, *Using the Bayesian Improved Surname Geocoding Method (BISG) to create a working classification of race and ethnicity in a diverse managed care population: a validation study*, Health Serv. Res. (2014), available at <https://pubmed.ncbi.nlm.nih.gov/23855558>. See also Rachad Alao, Miranda Bogen, Jingang Miao, Ilya Mironov, Jonathan Tannen, *How Meta is working to assess fairness in relation to race in the U.S. across its products and systems* (Nov. 2021), available at <https://ai.facebook.com/research/publications/how-meta-is-working-to-assess-fairness-in-relation-to-race-in-the-us-across-its-products-and-systems>.

³¹ See, e.g., Arthur P. Baines, Marsha J. Courchane, *Fair Lending: Implications for the Indirect Auto Finance Market* American Financial Services Association (2014), available at <https://media.crai.com/sites/default/files/publications/Fair-Lending-Implications-for-the-Indirect-Auto-Finance-Market.pdf>; Jiahao Chen, Nathan Kallus, Xiaojie Mao, Geoffry Svacha, Madeleine Udell, *Fairness Under Unawareness: Assessing Disparity When Protected Class Is Unobserved*, Conference on Fairness, Accountability, and Transparency (2019), available at <https://dl.acm.org/doi/10.1145/3287560.3287594>.

³² Airbnb Anti-Discrimination Team, *Measuring Discrepancies in Airbnb Guest Acceptance Rates Using Anonymized Demographic Data* (Jun. 12, 2020), available at <https://news.airbnb.com/wp-content/uploads/sites/4/2020/06/Project-Lighthouse-Airbnb-2020-06-12.pdf>.

³³ *Id.*, at 13.

³⁴ Airbnb, *A new way we’re fighting discrimination on Airbnb* (Jun. 15, 2020), <https://www.airbnb.com/resources/hosting-homes/a/a-new-way-were-fighting-discrimination-on-airbnb-201>.

³⁵ White House, *A Vision for Equitable Data: Recommendations from the Equitable Data Working Group* at 7 (Apr. 2022), available at <https://www.whitehouse.gov/wp-content/uploads/2022/04/eo13985-vision-for-equitable-data.pdf>.

When underlying law, policy, or regulation requires an entity to *not* collect race or ethnicity information, inference methodologies are one of the remaining tools available to measure and detect disparities. Given that, “[r]egulators should continue to research ways to further improve protected class status imputation methodologies using additional data sources and more advanced mathematical techniques.”³⁶ Beyond ongoing interagency research that could improve the accuracy of inference methodologies, agencies should also clearly consider interagency guidance that would offer companies a framework for how to use inference methodologies for anti-discrimination testing purposes. Such guidance could detail under what circumstances certain inference methods should be used and how to choose a methodology for inferring this data given its sensitivity and context-dependent nature. For example, “methods to generate data on perceived race may be relevant to assess the impact of race in people’s interactions with one another, but not so much for understanding whether a product or system performs differently across self-identified races and ethnicities.”³⁷ Indeed, “when measuring disparities, the choice of inference or racial proxy might be an important part of the analysis design. If one wants to measure disparities, one should have in mind what might be driving those disparities, or at least be willing to speculate and explore a range of potential causes.”³⁸

At the same time, it is worth explicitly re-considering existing prohibitions in civil rights law on the collection of demographic data. For example, the Federal Reserve Board (FRB) twice considered and ultimately rejected proposals to lift Regulation B’s prohibition on collection of demographic data for non-mortgage credit products in 1996 and 2003.³⁹ Notably, when the FRB reconsidered lifting the prohibition on the collection of demographic data for the second time, it did so in no small part “in response to concerns that continue to be expressed by the Department of Justice and some of the federal financial enforcement agencies, pointing to anecdotal evidence of discrimination in

³⁶ Michael Akinwumi, John Merrill, Lisa Rice, Kareem Saleh, Maureen Yap, *An AI fair lending policy agenda for the federal financial regulators*, Brookings Center on Regulation and Markets (Dec. 2021), available at <https://www.brookings.edu/research/an-ai-fair-lending-policy-agenda-for-the-federal-financial-regulators/>.

³⁷ Rachad Alao, Miranda Bogen, Jingang Miao, Ilya Mironov, Jonathan Tannen, *How Meta is working to assess fairness in relation to race in the U.S. across its products and systems* (Nov. 2021), available at <https://ai.facebook.com/research/publications/how-meta-is-working-to-assess-fairness-in-relation-to-race-in-the-us-across-its-products-and-systems>.

³⁸ Aaron Rieke, Vincent Southerland, Dan Svirsky, Mingwei Hsu, *Imperfect Inferences: A Practical Assessment*, Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (2022), available at <https://dl.acm.org/doi/fullHtml/10.1145/3531146.3533140>.

³⁹ Board of Governors of the Federal Reserve System, *Withdrawal of Proposed Rule re: Equal Credit Opportunity*, 61 FR 68688 (Dec. 30, 1996); Board of Governors of the Federal Reserve System, *Final Rule re: Equal Credit Opportunity*, 68 FR 13143 (Mar. 18, 2003).

connection with small business and other types of credit. These agencies believe that the ability to obtain and analyze data about race and ethnicity (such as creditors might collect on a voluntary basis) would aid fair lending enforcement.”⁴⁰ In particular, “most of the federal financial enforcement agencies, the Department of Justice, [and] the Department of Housing and Urban Development . . . favored removing the prohibition.”⁴¹

There was good reason that these fair lending regulators wanted the removal of Regulation B’s prohibition on collection of demographic data. As one letter from a Federal Reserve Bank to the FRB noted, “its examiners were unable to conduct thorough fair lending examinations or review consumer complaints alleging discrimination for nonmortgage products due to the lack of available data.”⁴² The FRB ultimately decided they would permit the collection of data on race, color, gender, national origin, religion, marital status, and age so long as this collection was in connection with a self-test, where the results of that test are privileged. However, the voluntary self-testing regime has proven to be ineffective, as “[v]anishingly few creditors take advantage of this exception.”⁴³

4. The federal government should provide clear guidance to companies about measuring algorithmic discrimination and searching for less discriminatory alternative models.

The NTIA should suggest a fundamental reorientation regarding the search for less discriminatory alternatives: companies in traditionally covered civil rights areas should have an affirmative obligation to prevent and remedy algorithmic discrimination based on their testing, and to search for less discriminatory alternative models.

The search for less discriminatory alternatives frequently takes place in two contexts. First, traditional fair lending testing and compliance requires lenders to assess whether their models lead to negative outcomes for protected classes. If they do, lenders are supposed to ensure that the model serves a legitimate business need and determine

⁴⁰ Board of Governors of the Federal Reserve System, *Proposed Rule re: Equal Credit Opportunity*, 64 FR 44582, 44585 (Aug. 16, 1999), available at <https://www.govinfo.gov/content/pkg/FR-1999-08-16/pdf/99-20598.pdf>.

⁴¹ *Id.*, at 44586.

⁴² United States Government Accountability Office, *Fair Lending: Race and Gender Data Are Limited for Nonmortgage Lending* at 18 (Jun. 2008), available at <https://www.gao.gov/assets/gao-08-698.pdf>.

⁴³ Rebecca Kelly Slaughter, Janice Kopec, Mohamad Batal, *Algorithms and Economic Justice: A Taxonomy of Harms and a Path Forward for the Federal Trade Commission*, 23 Yale J.L. & Tech 1, 43 (2021), available at https://law.yale.edu/sites/default/files/area/center/isp/documents/algorithms_and_economic_justice_master_final.pdf.

whether changes to the model would result in less of a disparate effect. While some financial institutions do routinely test their models for discrimination risks and search for less discriminatory alternatives, many do not.⁴⁴ Second, the search for less discriminatory alternatives sometimes takes place in the context of disparate impact litigation. But because the burden falls on the plaintiff challenging an allegedly discriminatory policy or practice to identify such an alternative, it's often infeasible.

Overall, there are few affirmative obligations for companies to explicitly search for less discriminatory alternative models. Meanwhile, voluntary efforts are often stymied by a longstanding (but misguided) assumption that a trade-off exists between a model's accuracy and the fairness of its outcomes.⁴⁵

Recent research has shown that it is often possible to develop many different, equally accurate algorithmic models that differ in the degree to which they result in disparities in outcomes across groups — even when using the same target, features, and training data.⁴⁶ That is, there is not necessarily only one accurate model for a given task, and there is not always a trade-off between a model's accuracy and a model's disparities.⁴⁷ As a recent paper explains, there are usually multiple models with equivalent accuracy, but significantly different properties.⁴⁸ This phenomenon is called “model multiplicity,” which describes when “models with equivalent accuracy for a certain prediction task differ in terms of their internals—which determine a model's decision process—and their predictions.”⁴⁹ This multiplicity “creates the possibility to minimize differences in prediction-based metrics across groups, notably differential validity (*i.e.*, differences in accuracy) and disparate impact (*i.e.*, differences in model predictions).”⁵⁰

⁴⁴ See Task Force on Artificial Intelligence Testimony of Stephen F. Hayes, *Equitable Algorithms: How Human-Centered AI Can Address Systemic Racism and Racial Justice in Housing and Financial Services*, 117th Cong. (May 7, 2021), available at <https://financialservices.house.gov/uploadedfiles/hhrg-117-ba00-wstate-hayess-20210507.pdf>.

⁴⁵ Maria De-Arteaga, Stefan Feuerriegel, Maytal Saar-Tsechansky, *Algorithmic Fairness in Business Analytics: Directions for Research and Practice*, *Production and Operations Management*, 31, 3749–3770 (2022).

⁴⁶ Charles T. Marx, Flávio P. Calmon, Berk Ustun, *Predictive Multiplicity in Classification*, In Proceedings of the 37th International Conference on Machine Learning (2020).

⁴⁷ Kit T. Rodolfa, Hemank Lamba, Rayid Ghani, *Empirical Observation of Negligible Fairness–Accuracy Trade-offs in Machine Learning for Public Policy*, 3 *Nat. Mach. Intel.* 896-904 (2021), available at <https://www.nature.com/articles/s42256-021-00396-x>.

⁴⁸ Emily Black, Manish Raghavan, Solon Barocas, *Model Multiplicity: Opportunities, Concerns, and Solutions*, Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency (2022), available at <https://dl.acm.org/doi/10.1145/3531146.3533149>.

⁴⁹ *Id.* at 850.

⁵⁰ *Id.* at 854.

Practically speaking, model multiplicity directly bears on the search for less discriminatory alternative models. Because of model multiplicity, both plaintiffs that bring discrimination lawsuits, as well as defendants who use algorithmic systems, would theoretically appear much better positioned to identify a less discriminatory alternative model than they have been in the past. However, “[n]either agencies nor courts have delineated concrete thresholds for determining whether a less discriminatory alternative practice must be adopted because it sufficiently achieves a legitimate business need.”⁵¹ Thirty years ago, a law review article noted that Congress and regulators had provided “scant guidance as to when a proposed [less discriminatory alternative] is sufficiently less discriminatory to warrant imposition on an employer.”⁵² Unfortunately, little has changed since. For example, does an alternative practice need to be “equally effective” to survive scrutiny? If so, what does it mean for a practice to be “equally effective”? Are costs relevant? If so, how? Each of these pressing questions do not have clear answers in law, regulation, or policy.

In fact, existing guidance offers somewhat conflicting answers to some of these questions. For example, under regulations implementing the Fair Housing Act, a plaintiff does not need to demonstrate that a less discriminatory alternative is “equally effective,” but instead that the alternative serves a defendant’s legitimate interests.⁵³ Under Title VII, the 1989 Supreme Court decision in *Wards Cove v. Atonio* followed by the Civil Rights Act of 1991 together left a muddled and imprecise standard for less discriminatory alternatives.⁵⁴

⁵¹ Relman Colfax PLLC, *Fair Lending Monitorship of Upstart Network’s Lending Model*, Third Public Report of the Independent Monitor at 13 (Sep. 16, 2021), available at https://www.relmanlaw.com/media/cases/1333_PUBLIC%20Upstart%20Monitorship%203rd%20Report%20FINAL.pdf.

⁵² Note, *The Civil Rights Act of 1991 and Less Discriminatory Alternatives in Disparate Impact Litigation*, 106 Harv. L. Rev 1621, 1627 (1993).

⁵³ Department of Housing and Urban Development, *Implementation of the Fair Housing Act’s Discriminatory Effects Standard*, 78 FR 11459 (Feb. 15, 2013), available at <https://www.federalregister.gov/documents/2013/02/15/2013-03375/implementation-of-the-fair-housing-acts-discriminatory-effects-standard>.

⁵⁴ The Civil Rights Act of 1991 abrogated the Supreme Court’s decision in *Wards Cove*, which held that a plaintiff’s proposed less discriminatory alternative must be “equally effective” as the defendant’s chosen policy at serving the defendant’s interest, taking into account factors such as the cost or other burdens that alternative policies would impose. *Wards Cove Packing Company, Inc., et al., v. Frank Atonio*, 490 U.S. 642 (1989). Under the Civil Rights Act of 1991, a plaintiff’s demonstration of an alternative employment practice “shall be in accordance with the law as it existed on June 4, 1989” — the day before the Court decided *Wards Cove*. 42 U.S.C. 2000e-2(k)(1)(C). Unfortunately, “the exact nature of the ‘less discriminatory alternative’ standard was far less than clear even before *Wards Cove*.” Matthew U. Scherer, Allan G. King & Marko J. Mrkonich, *Applying Old Rules to New Tools: Employment Discrimination Law in the Age of Algorithms*,

Outside of Title VII, the Supreme Court has nonetheless applied the *Wards Cove* articulation of less discriminatory alternatives to other anti-discrimination statutes.⁵⁵

As has been readily documented, in the traditionally covered civil rights areas such as credit, employment, and housing, companies routinely deploy algorithmic systems. These systems are often evaluated based on quantitative measures of accuracy, allowing precise comparison of their relative benefits. And in many cases, it will be possible to exploit model multiplicity to avoid trading off accuracy in favor of reducing disparate impact — that is, to find a clearly less discriminatory, but similarly accurate, alternative practice. As a result, agencies will need to provide more precise guidance on what makes a less discriminatory alternative sufficient to warrant imposition. That guidance should clearly address how similarly effective an alternative practice needs to be relative to an existing, challenged practice. It also should carefully scope which costs are relevant in determining if an alternative practice is actually viable.

5. The federal government should significantly expand its anti-discrimination testing capabilities to uncover algorithmic discrimination.

In several civil rights domains, demographic testing has been a historically important means of rooting out discrimination. The federal government has a long history of using covert testing to uncover evidence of discrimination by landlords, lenders, and others. Since “the late 1970s, the Department of Housing and Urban Development (HUD) has monitored the forms and incidence of racial and ethnic discrimination in both rental and sales markets approximately once a decade through nationwide paired-testing studies.”⁵⁶ Since its formation in 1991, the Department of Justice’s Fair Housing Testing Program has sought to uncover unlawful discrimination under statutes like the Fair

71 S.C. L.Rev. 449, 471, (2019). In fact, though the statute says the demonstration of an alternative employment practice shall be in accordance with the law as it existed on June 4, 1989, “the phrase ‘alternative employment practice’ is absent from every Supreme Court disparate impact decision preceding *Wards Cove*.” Michael Zimmer, *Individual Disparate Impact Law: On the Plain Meaning of the 1991 Civil Rights Act*, 30 Loy. U. Chi. L. J. 473, 485 (1999).

⁵⁵ *Hardie v. Nat’l Collegiate Athletic Ass’n*, 876 F.3d 312 (9th Cir. 2017) (noting that the Supreme Court has continued to apply *Wards Cove* burden shifting to other anti-discrimination statutes).

⁵⁶ Claudia L. Aranda, *Housing Discrimination in America: Lessons From the Last Decade of Paired-Testing Research*, Statement before the Subcommittee on Transportation, Housing and Urban Development and Related Agencies at 3 (Feb. 27, 2019), available at https://www.urban.org/sites/default/files/publication/99836/housing_discrimination_in_america_-_claudia_aranda.pdf.

Housing Act, Title II of the Civil Rights Act of 1964, the Equal Credit Opportunity Act, and the Americans with Disabilities Act. Such affirmative testing is critical “[w]here discrimination is hidden or hard to detect [and it] provides an indispensable tool for uncovering and exposing discriminatory policies and practices.”⁵⁷ Beyond DOJ and HUD, a few other agencies have also piloted testing programs from the Federal Reserve Board, to the Office of the Comptroller of the Currency, to the Equal Employment Opportunity Commission, and the Office of Federal Contract Compliance Programs.⁵⁸

Just as the federal government stood up anti-discrimination testing efforts to detect discrimination in the physical world, it must do so to detect discrimination in digital systems. The NTIA should advise the President that the relevant agencies that enforce existing civil rights laws need to develop federal standards for conducting civil rights audits and assessments of algorithmic systems that affect covered areas, such as housing, jobs, and lending. This includes developing new testing methods to uncover discrimination in digital systems that mediate access to economic opportunities, under the government’s own testing programs.

It also requires new resources and capacity: either hiring technologists who can help relevant civil rights agencies conduct civil rights audits or providing funding for external technologists to conduct the same. The NTIA should support the federal government to create the necessary programs and initiatives to bring anti-discrimination testing capacity into the government, and to create agency infrastructure as appropriate to integrate and embed such testing into agency work.

We welcome further conversations on these important issues. If you have any questions, please contact Logan Koepke (Project Director, logan@upturn.org) and Harlan Yu (Executive Director, harlan@upturn.org).

⁵⁷ Department of Justice, *The DOJ Fair Housing Testing Program: Three Decades of Guarding Civil Rights* at 4 (Apr. 2022), available at <https://www.justice.gov/crt/page/file/1497551/download>.

⁵⁸ Government Accounting Office, *Fair Lending: Federal Oversight and Enforcement Improved but Some Challenges Remain* at 61 (Aug. 1996), available at <https://www.govinfo.gov/content/pkg/GAOREPORTS-GGD-96-145/pdf/GAOREPORTS-GGD-96-145.pdf>; see also Darrick Hamilton, Rebecca Dixon, *Shifting the Burden of Proof: Using Audit Testing to Proactively Root Out Workplace Discrimination* (Sep. 2022), available at <https://www.nelp.org/publication/using-audit-testing-to-proactively-root-out-workplace-discrimination>.